Why geographic redundancy is essential for Cloud - blueAPACHE

Canonical: https://www.blueapache.com/blog/why-geographic-redundancy-is-essential-to-cloud-success

Details:

Historical article: 2016. Infrastructure and availability statements below describe the offering reported at that time; current commitments depend on the agreed service.

Geographic redundancy – no storm in a teacup

Earlier this month, Amazon Web Services' (AWS) experienced crippling power outage in their Sydney datacenters due to wild weather that swept across coastal New South Wales. Torrential rains and strong winds knocked out the operations of numerous AWS customers across Australia, preventing end users from accessing websites and other critical systems. According to reports, AWS clients including Domino's Pizza, Foxtel, The Iconic, Stan and Domain were unable to trade from mid-afternoon until as late as the next morning. Although the power was restored within about 90 minutes, it took hours for some of the systems to be rebooted and brought back online, leading to substantial losses. As more businesses move critical applications and systems to the cloud, the impact a natural disaster or an unexpected outage can have on your business operations is greater today. Even small outages can have long term business impacts. Redundant components within a single datacentre are important and readily available; but what if the datacentre itself is hit by a natural disaster? This is where geographic redundancy becomes important. It ensures high availability of business critical systems across multiple locations, mitigating the risk of weather outages as experienced by AWS. Businesses can mitigate downtime by replicating applications and data across multiple 'geo-diverse' locations. Also termed as 'geo-replication', the data that is created or updated in a primary location, is asynchronously replicated to a secondary location so that the same data exists and is readily accessible in both locations. Ideally the datacentre locations are geographically separated (Melbourne and Sydney for example), so that should one of the locations experience a catastrophic event and cannot be restored, the secondary location can quickly and seamlessly take over the primary role. All traffic is automatically rerouted to the secondary site with minimal service downtime for users. In order for businesses to ensure that critical data and applications remain highly available, geographical redundancy is critical. Our cloud environment is built on a world-class service provider platform that is located in three core geographically diverse datacentres along Australia's eastern seaboard. This redundancy enables us to provide a 99.999% uptime guarantee. To learn more about geographic redundancy, or business continuity and disaster recovery, contact our account team .

What incident prompted your blog post on geographic redundancy?

In June 2016, Amazon Web Services (AWS) experienced a crippling power outage in its Sydney datacentres caused by wild weather—torrential rains and strong winds—sweeping across coastal New South Wales, which knocked out operations for numerous AWS customers across Australia.

How long did it take for AWS's power to be restored, and why did the outage still cause substantial losses?

The power was restored within about 90 minutes, but it took hours for some of the affected systems to be rebooted and brought back online, leading to substantial losses for the businesses involved.

What is geographic redundancy according to the blueAPACHE article?

Geographic redundancy ensures high availability of business-critical systems across multiple locations, mitigating the risk of an entire datacentre being taken offline by events like weather outages, as experienced by AWS.

Where is our cloud environment's infrastructure located, and what uptime guarantee does it provide?

The 2016 article described three geographically diverse data centres along Australia's eastern seaboard and advertised 99.999% uptime. That is a historical description; confirm the infrastructure, applicable service level and exclusions for a current proposal.

How can businesses learn more about geographic redundancy or business continuity and disaster recovery from you?

We advise contacting our account team via our contact page to learn more about geographic redundancy, business continuity, and disaster recovery.

What should a customer confirm about STaaS resilience?

The historical STaaS account describes geographically distributed storage resilience. It does not establish an unconditional guarantee that every health record remains secure and accessible in real time after any facility compromise. Confirm replication design, recovery targets, supported failures and testing for the proposed service.

How does geographic redundancy relate to your disaster recovery service, emPOWER IT Continuity Services (DRaaS)?

Geographically separated infrastructure can support disaster-recovery design, but recovery depends on the systems, replication, failover procedures and testing in scope. Confirm Recovery Time Objective and Recovery Point Objective commitments in the applicable IT Continuity Services agreement.